arXiv:1603.02882v2 [math.PR] 8 Aug 2016

A Conjugate Approach to Partially Observable Markov Decision Processes on Borel Spaces

Yun Shen[†] Wilhelm Stannat[*] Klaus Obermayer[†]

July 22, 2016

Abstract

This paper presents a conjugate approach to solve the optimization problem of partially observable Markov decision processes (POMDPs) on Borel spaces with unbounded reward functions. We equip the belief state space of probability distributions with the Wasserstein metric and construct thereby a weighted norm for the functions on belief states. Under the weighted norm, we proof the existence of an optimal solution to POMDPs with possibly unbounded reward functions on both sides by reducing them to canonical Markov decision processes (MDPs). Furthermore, we present a conjugate duality theorem of the Fenchel-Moreau type on the Wasserstein space of belief states. Applying the conjugate duality to POMDPs, we derive an iterative conjugate solution in form of a convex and lower semicontinuous function by iterating the level sets, which is shown to be arbitrarily close to the optimal solution to POMDPs.

Key words: partially observable Markov decision processes; Wasserstein metric; conjugate duality; Fenchel-Moreau theorem; weighted norm; optimal control

MSC2000: 90C40, 90C46, 90C25, 46B10

1 Introduction

In many real-world problems, decisions have to be made in face of incomplete information. *Partially observable Markov decision processes* (POMDPs, Smallwood and Sondik (1973)) are proposed to deal with the scenarios where the underlying states are not fully observable. POMDPs have important applications in various fields, e.g., operations research (Lovejoy (1991)), robotics (Pineau et al. (2006)), artificial intelligence (Kaelbling et al. (1998)), and most recently, computational neuroscience (Rao (2010)).

It is well known that a POMDP can be reduced to a standard *Markov decision process* (MDP) where states are completely observable. We refer to Sondik (1978) for finite spaces, to Sawaragi and Yoshikawa (1970) for countable spaces, and to Hernández-Lerma (1989, Chapter 4) and Feinberg et al. (2016) for Borel spaces. For example, in the setting of finite state spaces with a discounted infinite-horizon objective, the optimal solution is given by solving the following Bellman equation

$$\phi(\mu) = \max_a \left\{ r(\mu, a) + \alpha \sum_y \phi(\mu'(\mu, a, y)) \left(\sum_{x'} P(x'|\mu, a) Q(y|x', a) \right) \right\}, \forall \mu \in \mathscr{P}. \quad (1)$$

Here, x', a and y denote the hidden (or latent) state, the action and the observation respectively, and $\mu \in \mathscr{P}$ is the distribution over states (also called a belief state), while

[*]Institut für Mathematik, and Bernstein Center for Computational Neuroscience, Technische Universität Berlin, Straße des 17. Juni 136, 10623, Berlin, Germany (stannat@math.tu-berlin.de).

[†]Fakultät Elektrotechnik und Informatik, and Bernstein Center for Computational Neuroscience, Technische Universität Berlin, Marchstr. 23, 10587, Berlin, Germany (yun.shen@tu-berlin.de, klaus.obermayer@mailbox.tu-berlin.de).

μ' is the posterior distribution of the successive state given by

$$\mu'(\cdot|\mu, a, y) = \frac{P(\cdot|\mu, a)Q(y|\cdot, a)}{\sum_{x'} P(x'|\mu, a)Q(y|x', a)}.$$

P is controls the transition probability between states, Q governs the observation probability of y given states and actions, r denotes the reward function, and $\alpha \in (0, 1)$ serves as a discount factor. We will state a formal introduction of POMDPs on Borel spaces in Section 2.1.

In spite of knowing this theoretical solution, until recently, POMDPs were notoriously difficult to solve in practice, even with a handful of hidden states (Shani et al. (2013)), since in the Bellman equation (1), the optimal "value function" ϕ has to be solved for all distributions $\mu \in \mathscr{P}$, which contains infinite number of points even for a finite hidden state space. Over the past decade, a significant breakthrough has been made in the development of POMDPs solving algorithms, largely due to an approach called *point-based value iteration* (Pineau et al. (2006)). It computes the value function ϕ over a finite subset of the distribution space with a controllable approximation error. The success of this algorithm is largely attributed to the property of ϕ observed by Smallwood and Sondik (1973) for POMDPs with finite spaces: The optimal solution to the Bellman equation (1) can be arbitrarily well approximated by a piecewise linear function $\hat{\phi}$ with following representation

$$\hat{\phi}(\mu) = \sup_{h \in N} \sum_{x} h(x)\mu(x) \tag{2}$$

where N is a finite set of functions on the hidden state space. Applying this representation to the Bellman equation, the sum on the right-hand side of (1) becomes the form

$$\sum_{y} \left(\sup_{h \in N} \sum_{x'} h(x')P(x'|\mu, a)Q(y|x', a) \right),$$

which is the key to designing fast computing algorithms.

To our best knowledge, it is still an open question whether the piecewise linear approximation discovered by Smallwood and Sondik (1973) for finite-space POMDPs holds also for POMDPs on general spaces. The major contribution of this paper is, therefore, to provide a confirmative answer to this open question, by applying a conjugate approach to the space of probability measures, equipped with the Wasserstein metric.

Our first intuition is that the representation stated in (2) is in fact of the Fenchel conjugate form, if ϕ is a convex and lower semicontinuous function on $\mathscr{P}$. However, the standard Fenchel-Moreau theorem (see e.g. Ioan-Bot et al. (2009, Chapter 2)) in the convex analysis cannot be directly applied here, since it is difficult to identify dual spaces of the probability measures $\mathscr{P}$ on general spaces, with the usual topologies induced by, e.g., the metric of weak convergence or the metric of total variation. To overcome this difficulty, our trick is to apply the Wasserstein metric of order 1 (see e.g. Villani (2009), also known as Kontrovich-Rubinstein distance). Endowed with the induced topology, we obtain a Fenchel-Moreau-type dual representation with Lipschitz functions. That is, the set N in (2) is a subset of Lipschitz functions on the hidden state space. Thereby we show that it allows a dual representation similar to (2).

Secondly, we will show that the convexity and lower semicontinuity of the value function ϕ remain unchanged in the course of value iteration, where we iterate ϕ as on the right-hand side of (1). This confirms the existence of an arbitrarily good approximation of the optimal value function with a convex and lower semicontinuous function that allows a dual presentation as in (2).

Finally, since the topology, with which the probability measure space is endowed, differs from those applied in POMDP-literature, we adapt the whole framework accordingly. We construct the weighted norm space of functions on the Wasserstein space. Under mild assumptions, we show that POMDPs with possibly unbounded reward functions on both sides can be reduced to standard MDPs equipped with weighted norm, where the existence of an optimal solution is guaranteed. This leads to another contribution of this paper: Our results extend the domain of potential applications, in comparison with the literature of POMDPs on Borel spaces (see Hernández-Lerma (1989); Feinberg et al. (2016)), where the reward function has to be at least one-side bounded.

The paper is organized as follows. In Section 2, we introduce the setup of discounted POMDPs on Borel spaces and the concept of the Wasserstein metric, as well as its basic properties. In Section 3, POMDPs are reduced to MDPs on Borel spaces of probability measures equipped with a weight norm that is specifically constructed with the metric of the hidden state space. Adapting to an MDP with the weighted norm, we show the existence of an optimal solution with an optimal deterministic policy to the POMDP. In Section 4, we present a duality theorem of the Fenchel-Moreau type on Wasserstein space. The duality is then applied to POMDPs and we obtain that under mild assumptions, the optimal solution to the POMDP can be arbitrarily approximated by a convex and lower semicontinuous function of the dual form. Finally, in Section 5, we present an example of POMDP with dynamics given by a 1-dimensional nonlinear Kalman filter, to demonstrate the applicability of the developed theory.

Notation A *Polish space* is a complete separable metric space and a *Borel space* is a Borel subset of a Polish space. If X is a Borel space, its Borel σ-algebra is denoted by $\mathcal{B}(\mathsf{X})$. Let X and Y be two Borel spaces. A *stochastic kernel on* X *given* Y is a function $\psi(B|y), B \in \mathcal{B}(\mathsf{X}), y \in \mathsf{Y}$ such that i) $\psi(\cdot|y)$ is a probability measure on $\mathcal{B}(\mathsf{X})$ for every fixed $y \in \mathsf{Y}$, and ii) $\psi(B|\cdot)$ is a measurable function on Y for every fixed $B \in \mathcal{B}(\mathsf{X})$.

2 Model description and setup

2.1 POMDPs

A partially observable Markov decision process (POMDP, see e.g., Hernández-Lerma (1989, Chapter 4)) is composed of $(\mathsf{X}, \mathsf{Y}, \mathsf{A}, P, Q, \mu, r)$, where:

(a) X, the (hidden or latent) state space, is a Polish space. Let its metric be denoted by d.

(b) Y, the space of observation, is a Borel space.

(c) A, the action space, is a Borel space.

(d) $P(dx'|x,a)$, the state transition law, is a stochastic kernel on X given $\mathsf{K} := \mathsf{X} \times \mathsf{A}$. K is also a Borel space.

(e) $Q(dy|a,x)$, the observation kernel, is a stochastic kernel on Y given K. Q_0, the initial observation kernel, is a stochastic kernel on Y given X.

(f) $\mu \in \mathscr{P}$ is the initial distribution with $\mathscr{P}$ being the space of all probability measures on $(\mathsf{X}, \mathcal{B}(\mathsf{X}))$.

(g) $r : \mathsf{K} \to \mathbb{R}$ is the one-step reward function, which is $\mathcal{B}(\mathsf{K})$-measurable.

The POMDP evolves as follows. At time $t = 0$, the initial (hidden or latent) state x_0 follows a given prior distribution μ, while the initial observation y_0 is generated according to an initial observation kernel $Q_0(\cdot|x_0)$. If at time t the state of the system is x_t and the control $a_t \in \mathsf{A}$ is applied, then the agent receives a reward $r(x_t, a_t)$ and the system transits to state x_{t+1} according to the transition law $P(dx_{t+1}|x_t, a_t)$. The observation y_{t+1} is generated by the observation kernel $Q(dy_{t+1}|a_t, x_{t+1})$. The *observed* history is defined as

$$h_0 := \{\mu, y_0\} \in \mathsf{H}_0 \text{ and } h_t := \{\mu, y_0, a_0, \ldots, y_{t-1}, a_{t-1}, y_t\} \in \mathsf{H}_t, t = 1, 2, \ldots, \quad (3)$$

where $\mathsf{H}_0 := \mathscr{P} \times \mathsf{Y}$ and $\mathsf{H}_{t+1} = \mathsf{H}_t \times \mathsf{Y} \times \mathsf{A}, t = 0, 1, \ldots$. Notably, comparing with the canonical Markov decision processes (MDPs, see e.g., Hernández-Lerma (1989)), the states $\{x_t\}$ are not observable and hence, a *policy* depends only on the observed history.

A deterministic policy $\boldsymbol{\pi} := [\pi_0, \pi_1, \ldots]$ is composed of a sequence of one-step policy $\pi_t : \mathsf{H}_t \to \mathsf{A}$, given the observed history up to time t. Let Π be the set of all deterministic policies. Note that even with extension to nondeterministic policies, it is known (Hernández-Lerma (1989, Chapter 4)) that an optimal policy to POMDP is always deterministic. Hence, we consider in this paper only deterministic policies. The Ionescu-Tulcea theorem (see e.g., Bertsekas and Shreve (1978, pp. 140–141)) implies that for each $\pi \in \Pi$ and an initial $\mu \in \mathscr{P}$, along with P, Q and Q_0, a probability measure $\mathbb{P}_\mu^{\boldsymbol{\pi}}$ and a stochastic process $\{X_t, Y_t, A_t\}$ can be defined on $(\mathsf{X} \times \mathsf{Y} \times \mathsf{A})^\infty$ in a canonical way endowed with the

σ-algebra defined by the products of Borel σ-algebras $\mathcal{B}(\mathsf{X})$, $\mathcal{B}(\mathsf{Y})$, and $\mathcal{B}(\mathsf{A})$. We denote by $\mathbb{E}_\mu^{\boldsymbol{\pi}}$ the expectation with respect to (w.r.t.) this probability measure $\mathbb{P}_\mu^{\boldsymbol{\pi}}$.

We consider the following discounted cumulative rewards

$$J_T(\boldsymbol{\pi}, \mu) := \mathbb{E}_\mu^{\boldsymbol{\pi}} \left[\sum_{t=0}^{T} \alpha^t r(X_t, A_t) \right], \tag{4}$$

where $\alpha \in (0,1)$ stands for a discount factor. The objective is then to maximize the expected reward over the set of deterministic policies Π,

$$\phi_T^*(\mu) := \sup_{\pi \in \Pi} J_T(\boldsymbol{\pi}, \mu), \quad \mu \in \mathscr{P}(\mathsf{X}).$$

When $T = \infty$, we write J_∞ and ϕ_∞^* as J and ϕ^* respectively.

2.2 Reduction to Markov decision process

We show briefly in this subsection that the POMDP can be reduced to a Markov decision process (MDP, see e.g. Hernández-Lerma and Lasserre (1999)). We follow mostly the derivation by Hernández-Lerma (1989, Chapter 4). We first introduce the following notation:

$$\tilde{r}(\mu, a) := \int r(x, a)\mu(dx), \quad \text{and} \quad \tilde{P}(\mathsf{B}|\mu, a) := \int \mu(dx) P(\mathsf{B}|x, a), \mathsf{B} \in \mathcal{B}(\mathsf{X}). \tag{5}$$

For any $\mathsf{C} \in \mathcal{B}(\mathsf{Y})$ and $\mathsf{B} \in \mathcal{B}(\mathsf{X})$, we define

$$R(\mathsf{B}, \mathsf{C}|\mu, a) := \int_{\mathsf{B}} Q(\mathsf{C}|a, x')\tilde{P}(dx'|\mu, a), \quad \text{and} \quad \tilde{R}(\mathsf{C}|\mu, a) := R(\mathsf{X}, \mathsf{C}|\mu, a). \tag{6}$$

Proposition 2.1. *There exists a stochastic kernel M from $\mathscr{P} \times \mathsf{A} \times \mathsf{Y}$ to X such that for each $\mu \in \mathscr{P}$, $a \in \mathsf{A}, \mathsf{B} \in \mathcal{B}(\mathsf{X})$ and $\mathsf{C} \in \mathcal{B}(\mathsf{Y})$,*

$$R(\mathsf{B}, \mathsf{C}|\mu, a) = \int_{\mathsf{C}} M(\mathsf{B}|\mu, a, y)\tilde{R}(dy|\mu, a). \tag{7}$$

Proof. Direct application of Bertsekas and Shreve (1978, Corollary 7.27.1). □

Here, M can also be viewed as a mapping $\mathscr{P} \times \mathsf{A} \times \mathsf{Y} \to \mathscr{P}$. With a slight abuse of notation, let $M(\mu, a, y) := M(\cdot|\mu, a, y) \in \mathscr{P}$. Then the following stochastic kernel

$$\tilde{Q}(\mathsf{D}|\mu, a) := \int \mathbf{1}_{\mathsf{D}}(M(\mu, a, y))\tilde{R}(dy|\mu, a), D \subset \mathscr{P} \tag{8}$$

presents the new transition kernel on $\mathscr{P}$, whose metric will be specified in the next subsection.

At time t let $\mu_t \in \mathscr{P}$ be the posterior distribution. Then given an action $a_t \in \mathsf{A}$ and an observation $y_{t+1} \in \mathsf{Y}$, the *successive* distribution $\mu_{t+1} \in \mathscr{P}$ is exactly given by

$$\mu_{t+1} = M(\mu_t, a_t, y_{t+1}). \tag{9}$$

Note that μ_{t+1} is a random measure since y_{t+1} is a random variable with the distribution $\tilde{R}(\cdot|\mu_t, a_t)$. Hence, the POMDP can be reduced to a Markov decision process (MDP) with the *belief* state space $\mathscr{P}$, the action space A, the reward function $\tilde{r}$ on $\mathscr{P} \times \mathsf{A}$ and the transition kernel on belief states $\tilde{Q}$ defined above.

Let $\tilde{h}_t$ be a t-stage history for the MDP described above:

$$\tilde{h}_t := \{\mu_0, a_0, \ldots, \mu_{t-1}, a_{t-1}, \mu_t\} \in \tilde{\mathsf{H}}_t := \mathscr{P} \times (\mathsf{A} \times \mathsf{P})^t, t = 0, 1, \ldots,$$

where μ_0 is the initial posterior distribution and $\mu_t \in \mathscr{P}$ are recursively defined by (9). Given the original t-stage history h_t defined in (3), let $m_t : \mathsf{H}_t \to \tilde{\mathsf{H}}_t$ be the mapping such that $m_t(h_t) = \tilde{h}_t, \forall h_t \in \mathsf{H}_t, t = 0, 1, \ldots$. For a history-dependent MDP-policy $\boldsymbol{\delta} := [\delta_0, \delta_1, \ldots]$, where $\delta_t : \tilde{\mathsf{H}}_t \to \mathsf{A}$, we define its counterpart POMDP-policy as $\boldsymbol{\pi}^\delta = [\pi_0^\delta, \pi_1^\delta, \ldots]$, with $\pi_t^\delta(h_t) = \delta_t(m_t(h_t))$. Hence, if a policy $\boldsymbol{\delta}$ is optimal for the MDP, then its counterpart policy $\boldsymbol{\pi}^\delta$ is also optimal for the POMDP. For more details, we refer to Hernández-Lerma (1989, Chapter 4) and references therein.

After this reduction to MDP, it is well known that under proper assumptions (see, e.g., (Hernández-Lerma, 1989, Chapter 4) and Feinberg et al. (2016)), the optimal ϕ^* for the infinite-stage case satisfy the following *optimality equation*:

$$\phi(\mu) = \mathcal{T}(\phi)(\mu) := \sup_{a\in\mathsf{A}} \left(r(\mu,a) + \alpha \int \phi\left(M(\mu,a,y)\right) \tilde{R}(dy|\mu,a) \right), \forall \mu \in \mathscr{P}, \tag{10}$$

where $\mathcal{T}$ is an operator on the space of functions on probability measures in $\mathscr{P}$. We will show in Section 3 that under general assumptions allowing reward functions unbounded below and above, the existence of a solution of the above equation is guaranteed, as well as the existence of an optimal deterministic policy.

2.3 Wasserstein metric of order 1

In the literature, different topologies on $\mathscr{P}$ induced by different types of metric have been considered, such as weak convergence and total variation convergence (see, e.g., Feinberg et al. (2016)). In this paper, we will apply a new type of topology that is induced by the Wasserstein metric of order 1. Let (X, d) be a Polish space with metric d. Then the Wasserstein space of order 1 is defined as

$$\mathscr{P}_{W_1} := \left\{ \mu \in \mathscr{P} \;\middle|\; \int d(x_0,x)\mu(dx) < \infty \right\}$$

where $x_0 \in \mathsf{X}$ is arbitrary and its distance (called *Wasserstein metric*) is defined as

$$W_1(\mu,\nu) := \inf\left\{ \mathbb{E}\left[d(X,Y)\right] \mid \mathrm{law}(X) = \mu, \mathrm{law}(Y) = \nu \right\}. \tag{11}$$

For any real-valued function $f : \mathsf{X} \to \mathbb{R}$, its Lipschitz seminorm is defined as

$$\|f\|_{\mathrm{Lip}} := \sup_{x\neq y} \frac{|f(x) - f(y)|}{d(x,y)}.$$

Denote by $\mathscr{L} := \{f : \mathsf{X} \to \mathbb{R} | \|f\|_{\mathrm{Lip}} < \infty\}$ the space of all real-valued Lipschitz functions on X and by $[\mathscr{L}]_1 := \{f \in \mathscr{L} | \|f\|_{\mathrm{Lip}} \leq 1\}$ the unit ball of $\mathscr{L}$. Then, it can be shown that W_1 has the following dual representation

$$W_1(\mu,\nu) = \sup_{f\in[\mathscr{L}]_1} \left(\int f d\mu - \int f d\nu \right). \tag{12}$$

It is worth to be mentioned that W_1 is also called *Kantorovich-Rubinstein distance* in this dual form. For more detailed discussion, we refer to Villani (2009, Chapter 5). We have the following property (see (Villani, 2009, Theorem 6.18)): *If* (X, d) *is Polish,* $(\mathscr{P}_{W_1}, W_1)$ *is also Polish.*

The *weighted norm* is proved to be very useful when dealing with MDPs (see e.g. Hernández-Lerma and Lasserre (1999)). We now state the weighted norm induced by the Wasserstein metric. First, we specify a weight function $w : \mathsf{X} \to [1, \infty)$ as follows

$$w(x) := 1 + k \cdot d(x_0, x), x \in \mathsf{X}, \tag{13}$$

with some fixed $x_0 \in \mathsf{X}$ and positive constant $k > 0$. This definition implies $w(x_0) = 1$. One can easily verify that w is Lipschitz, and therefore continuous and measurable. Denote by $\mathscr{L}_w$ the space of all continuous functions on X such that

$$\|f\|_w := \sup_{x\in\mathsf{X}} \frac{|f(x)|}{w(x)} < \infty.$$

Let $\mathscr{P}_w \subset \mathscr{P}$ be the space of probability measures on X, μ, satisfying $\int w d\mu < \infty$, equipped with the following concept of weak convergence.

Definition 2.2. *(a)* μ_n *is said to converge weakly to* μ *in* $\mathscr{P}_w$*, if (i) for any bounded continuous function* f *on* X*, it follows that* $\int_{\mathsf{X}} f(x)\mu_n(dx) \to \int_{\mathsf{X}} f(x)\mu(dx)$ *as* $n \to \infty$*; and (ii)* $\int_{\mathsf{X}} w(x)\mu_n(dx) \to \int_{\mathsf{X}} w(x)\mu(dx)$ *as* $n \to \infty$*. (b) A function* $\phi : \mathscr{P}_w \to \mathbb{R}$ *is said to be lower (resp. upper) semicontinuous if*

$$\liminf_{n\to\infty} \phi(\mu_n) \geq \phi(\mu) \ (resp.\ \limsup_{n\to\infty} \phi(\mu_n) \leq \phi(\mu)),$$

whenever μ_n *converges weakly to* μ*.* ϕ *is said to be continuous, if* ϕ *is both lower and upper continuous.*

Remark 2.3. It is worth to be mentioned that the weak convergence defined above is stronger than the usual weak convergence, which requires only (i). Hence, to emphasize this difference, we call the latter *canonical weak convergence* throughout the rest of this paper.

PROPOSITION 2.4. *The following two statements are equivalent: (i) μ_n converges weakly to μ in $\mathscr{P}_w$ and (ii) $\int f d\mu_n \to \int f d\mu, \forall f \in \mathscr{L}_w$.*

PROOF. See (Villani, 2009, Definition 6.8). □

PROPOSITION 2.5. *Let w be defined as in (13) with some $k > 0$ and $x_0 \in \mathsf{X}$. Then (i) $\mathscr{P}_{W_1} = \mathscr{P}_w$ and (ii) the weak convergence defined in Definition 2.2 is equivalent to the convergence in $(\mathscr{P}_{W_1}, W_1)$, in other words, W_1 metrizes the weak convergence.*

PROOF. (i) is obvious by the definition. (ii) is a direct result of Theorem 6.9 of Villani (2009). □

In the rest of this paper, we always assume that the weight function satisfies (13) and hence $\mathscr{P}_w$ are used interchangeably with $\mathscr{P}_{W_1}$.

3 Optimal solution on weighted space

In this section, we shall reduce POMDPs to MDPs on Borel spaces equipped with a proper weighted norm, where we mainly follow the approach by Hernández-Lerma and Lasserre (1999, Chapter 8). Since the setting in this paper is slightly different from the one in Hernández-Lerma and Lasserre (1999), we restate an adapted version of the main results in Appendix A for the readers' convenience.

First of all, by using the Wasserstein space $\mathscr{P}_w = \mathscr{P}_{W_1}$, the belief state space is Polish, and therefore, a Borel space. We next specify the weight function and its weighted norm on $\mathscr{P}_w$. Let $w : \mathsf{X} \to [1, \infty)$ be the weight function as defined in (13). Define $\tilde{w} : \mathscr{P}_w \to [1, \infty)$ as

$$\tilde{w}(\mu) := \int w d\mu. \tag{14}$$

It is easy to check that $\tilde{w}$ is a continuous function and hence measurable on $(\mathscr{P}_w, W_1)$. Define the following space of functions on $\mathscr{P}_w$ with bounded $\tilde{w}$-norm:

$$\mathscr{B}_{\tilde{w}} := \left\{ \phi : \mathscr{P}_w \to \mathbb{R} \;\middle|\; \phi \text{ is } \mathcal{B}(\mathscr{P}_w)\text{-measurable}, \|\phi\|_{\tilde{w}} := \sup_{\mu \in \mathscr{P}_w} \frac{|\phi(\mu)|}{\tilde{w}(\mu)} < \infty \right\}.$$

In the next subsection, we shall specify assumption on the original POMDP to ensure the Assumptions A.1 – A.3 for MDPs.

3.1 Assumptions

We introduce the following assumption for the reward function.

Assumption 3.1. (i) There exists a positive constant $\bar{r} > 0$ such that $|r(x, a)| \leq \bar{r} w(x)$, for each $(x, a) \in \mathsf{K}$. (ii) For each $x \in \mathsf{X}$, $a \mapsto r(x, a)$ is upper semicontinuous.

PROPOSITION 3.2. *Under Assumption 3.1, (i) $|\tilde{r}(\mu, a)| \leq \bar{r}\tilde{w}(\mu), \forall (\mu, a) \in \mathscr{P}_w \times \mathsf{A}$; (ii) for each $\mu \in \mathscr{P}_w$, $a \mapsto \tilde{r}(\mu, a)$ is upper semicontinuous; (iii) for each $a \in \mathsf{A}$, $\mu \mapsto \tilde{r}(\mu, a)$ is continuous in $\mathscr{P}_w$.*

PROOF. (i) For each $(\mu, a) \in \mathscr{P}_w \times \mathsf{A}$, we have $|r(\mu, a)| \leq \int |r(x, a)| \mu(dx) \leq \bar{r} \int w d\mu = \bar{r}\tilde{w}(\mu)$.

(ii) Let $\{a_n, n = 1, 2, \ldots\}$ be a sequence of actions converging to a_0 and set $r_n(x) := r(x, a_n), n \in \mathbb{N}$. By Assumption 3.1(i), $r_n \leq \bar{r} w$. Applying the reversed Fatou's lemma, we obtain

$$\limsup_{n \to \infty} \int r_n d\mu \leq \int \limsup_{n \to \infty} r_n d\mu \leq \int r_0 d\mu,$$

where the last inequality is due to Assumption 3.1(ii).

Finally, (iii) is a direct result of Proposition 2.4(ii). □

Similar to the assumptions made in the literature of MDPs (see, e.g., (Hernández-Lerma and Lasserre, 1999, Assumptions 8.3.2 and 8.3.3)), we introduce the following assumption on the transition kernel P.

Assumption 3.3. (i) There exists a constant $\beta \in (0, \alpha^{-1})$ such that

$$\int w(x')P(dx'|x,a) \leq \beta w(x), \forall (x,a) \in \mathsf{X} \times \mathsf{A}.$$

(ii) For each $x \in \mathsf{X}$, $a \mapsto \int w(x')P(dx'|x,a)$ is continuous.

Under the above assumption, we show that the new probability measure $M(\mu, a, y)$ belongs to $\mathscr{P}_w$ in almost sure sense and furthermore, the transition kernel $\tilde{Q}$ satisfies that $\int \tilde{w}(\mu')\tilde{Q}(d\mu'|\mu, a)$ is continuous in a.

PROPOSITION 3.4. *Suppose Assumption 3.3 holds. Then for each $\mu \in \mathscr{P}_w$ and $a \in \mathsf{A}$, (i) $M(\mu, a, y) \in \mathscr{P}_w$, $\tilde{R}(\cdot|\mu, a)$-almost surely; (ii) $\int \tilde{w}(\mu')\tilde{Q}(d\mu'|\mu, a) \leq \beta\tilde{w}(\mu)$; and (iii) for each $\mu \in \mathscr{P}_w$, $a \mapsto \int \tilde{w}(\mu')\tilde{Q}(d\mu'|\mu, a)$ is continuous.*

PROOF. Fix arbitrary $(\mu, a) \in \mathscr{P}_w \times \mathsf{A}$. (i) Let $\mathsf{C} \in \mathcal{B}(\mathsf{Y})$ be a subset such that $\tilde{R}(\mathsf{C}|\mu, a) > 0$. Then we have

$$\begin{aligned}\int_{\mathsf{C}} \int_{\mathsf{X}} w(x')M(dx'|\mu, a, y)\tilde{R}(dy|\mu, a) &= \int_{\mathsf{X}} Q(\mathsf{C}|a, x')w(x')\tilde{P}(dx'|\mu, a) \\ &\leq \int_{\mathsf{X}} Q(\mathsf{Y}|a, x')w(x')\tilde{P}(dx'|\mu, a) \\ &\leq \int_{\mathsf{X}} w(x')\tilde{P}(dx'|\mu, a) \leq \beta \int w d\mu < \infty.\end{aligned}$$

It implies that $\int_{\mathsf{X}} w(x')M(dx'|\mu, a, y) < \infty$, $\tilde{R}(\cdot|\mu, a)$-almost surely, and hence (i) holds.

(ii) By definition, we have

$$\begin{aligned}\int \tilde{w}(\mu')\tilde{Q}(d\mu'|\mu, a) &= \int_{\mathsf{Y}} \tilde{w}(M(\mu, a, y))\tilde{R}(y|\mu, a) \\ \text{(by (14))} \quad &= \int_{\mathsf{Y}} \int_{\mathsf{X}} w(x')M(dx'|\mu, a, y)\tilde{R}(y|\mu, a) \\ \text{(by (7) and Fubini's theorem)} \quad &= \int_{\mathsf{X}} Q(\mathsf{Y}|a, x')w(x')\tilde{P}(dx'|\mu, a) \\ &\leq \beta \int w d\mu = \beta\tilde{w}(\mu).\end{aligned}$$

(iii) Note that the above calculation yields

$$\int \tilde{w}(\mu')\tilde{Q}(d\mu'|\mu, a) = \int w(x')\tilde{P}(dx'|\mu, a) = \iint_{\mathsf{X}\times\mathsf{X}} w(x')P(dx'|x, a)\mu(dx).$$

Let $\{a_n, n = 1, 2, \ldots\}$ be a sequence of actions converging to a_0 and set

$$f_n(x) := \int w(x')P(dx'|x, a_n), n \in \mathbb{N}.$$

Hence, the required continuity is equivalent to showing that $\lim_{n\to\infty} \int f_n d\mu = \int f_0 d\mu$. Indeed, by Assumption 3.3(i), we have $f_n \leq \beta w, \forall n \in \mathbb{N}$. The reversed Fatou's lemma implies $\limsup_{n\to\infty} \int f_n d\mu \leq \int \limsup_{n\to\infty} f_n d\mu = \int f_0 d\mu$. On the other hand, we have $f_n \geq -\beta w, \forall n \in \mathbb{N}$. Then the extended Fatou's lemma implies $\liminf_{n\to\infty} \int f_n d\mu \geq \int \liminf_{n\to\infty} f_n d\mu = \int f_0 d\mu$. Combining the above two inequalities yields the convergence. □

Assumption 3.5. For each $\mu \in \mathscr{P}_w$, there exist stochastic kernels M on $\mathscr{P}$ given $\mathscr{P}_w \times \mathsf{A} \times \mathsf{Y}$ and $\tilde{R}$ on Y given $\mathscr{P}_w \times \mathsf{A}$ satisfying (7) such that, if $\{a_n \in \mathsf{A}, n = 1, 2, \ldots\}$ converges to $a_0 \in \mathsf{A}$ as $n \to \infty$,

(i) there exists a subsequence $\{a_{n_k}\} \subset \{a_n\}$ and a measurable set $\bar{\mathsf{C}} \in \mathcal{B}(\mathsf{Y})$ such that $\tilde{R}(\bar{\mathsf{C}}|\mu, a_0) = 1$ and for all $y \in \bar{\mathsf{C}}$, $M(\mu, a_{n_k}, y)$ converges canonically weakly to $M(\mu, a_0, y)$;

(ii) for each $\mathsf{C} \in \mathcal{B}(\mathsf{Y})$, $\tilde{R}(\mathsf{C}|\mu, a_n) \to \tilde{R}(\mathsf{C}|\mu, a_0)$ as $n \to \infty$.

This assumption is inspired by Condition (c) in (Feinberg et al., 2016, Theorem 3.2). A sufficient condition for it will be discussed in the next section (see Remark 4.7).

PROPOSITION 3.6. *Under Assumptions 3.5 and 3.3, for each $\mu \in \mathscr{P}_w$, $a \mapsto \tilde{Q}(\cdot|\mu, a)$ is canonically weakly continuous.*

PROOF. This canonical weak continuity is guaranteed by Feinberg et al. (2016, Theorem 3.4). □

Finally, to guarantee the existence of one "selector", we assume

Assumption 3.7. A is compact.

Recall that the operator $\mathcal{T} : \mathscr{B}_{\tilde{w}} \to \mathscr{B}_{\tilde{w}}$ is defined as

$$\mathcal{T}_a(\phi)(\mu) := \tilde{r}(\mu, a) + \alpha \int \phi\left(M(\mu, a, y)\right) \tilde{R}(dy|\mu, a) \quad \text{and} \quad \mathcal{T}(\phi)(\mu) := \sup_{a \in \mathsf{A}} \mathcal{T}_a(\phi)(\mu).$$

Under Assumptions 3.1 – 3.5, it is guaranteed that for each $\mu \in \mathscr{P}_w$, $a \mapsto \mathcal{T}_a(\phi)(\mu)$ is upper-semicontinuous (for a proof, see (Hernández-Lerma and Lasserre, 1999, Lemma 8.3.7(a))). Hence, under the additional Assumption 3.7, the optimal a in the above optimization problem is always attainable in A (see, e.g., (Hernández-Lerma and Lasserre, 1999, Lemma 8.3.8(a))). Hence, from now on, we replace "sup" with "max".

3.2 Value iteration

The following *value iteration* is a widely used method to compute the optimal solution for POMDPs, and MDPs as well. Starting from arbitrary *value function* in $\mathscr{B}_{\tilde{w}}$, ϕ_0, at time t, we update value function as follows

$$\phi_{t+1}(\mu) = \mathcal{T}(\phi_t)(\mu) = \max_{a \in \mathsf{A}} \left(\tilde{r}(\mu, a) + \alpha \int \phi\left(M(\mu, a, y)\right) \tilde{R}(dy|\mu, a) \right)$$

Finally, by an adapted version of Theorem 8.3.6 by Hernández-Lerma and Lasserre (1999), which is stated as Theorem A.4 in Appendix A, we obtain the following convergence.

THEOREM 3.8. *Suppose that Assumptions 3.1 – 3.7 hold. Let β be the constant in Assumption 3.3(i) and $\bar{r}$ be the constant in 3.1(i) and define $\gamma := \alpha\beta \in (0, 1)$. Then*

(a) the optimal value function ϕ^* is the unique fixed point of the operator $\mathcal{T}$ satisfying $\phi^* = \mathcal{T}(\phi^*)$ in $\mathscr{B}_{\tilde{w}}$ and $\|\phi_t - \phi^*\|_{\tilde{w}} \leq \bar{r}\gamma^t/(1-\gamma), t = 1, 2, \ldots$.

(b) there exists a selector $f^* : \mathscr{P}_w \to \mathsf{A}$ such that

$$\phi^*(\mu) = \tilde{r}(\mu, f^*(\mu)) + \alpha \int \phi\left(M(\mu, f^*(\mu), y)\right) \tilde{R}(dy|\mu, f^*(\mu)), \forall \mu \in \mathscr{P}_w.$$

and $\boldsymbol{\pi}^* = (f^*)^\infty$ is one optimal policy satisfying $\phi^*(\mu) = J(\mu, \boldsymbol{\pi}^*), \forall \mu \in \mathscr{P}_w$.

PROOF. The original POMDP specified in Subsection 2.1 can be reduced to an MDP with $(\mathscr{P}_w, \mathsf{A}, \tilde{r}, \tilde{Q})$. Under Assumptions 3.1 – 3.7, Propositions 3.2 – 3.6 hold, and therefore, the conditions required by Theorem A.4 in Appendix A are satisfied. The assertion is then a direct application of Theorem A.4. □

Comparison with literature. Most early literature on POMDPs (see e.g. Sondik (1978); Hernández-Lerma (1989)) considers finite spaces or bounded reward functions only. In a recent work, Feinberg et al. (2016) consider a more general setting with Borel spaces and reward/cost functions that are bounded on one side, whereas the other side is allowed to be infinite. In this paper, we adopt the setting of MDPs in Hernández-Lerma and Lasserre (1999) with Borel spaces and real-valued reward functions. By introducing a weight function, the reward function can be possibly two-sided unbounded. Hence, we can cover many applications (??) that are not allowed in Feinberg et al. (2016). We would like also to point out that essentially, the proof in Feinberg et al. (2016) is based on the monotonicity argument, while our proof is obtained by applying the convergence under weighted norm.

4 A conjugate approach

4.1 A conjugate duality on Wasserstein space

Let (X, d) be a Polish space and $\mathscr{P}_{W_1}$ be the Wasserstein space of probability measures on X with order 1, equipped with the Wasserstein metric W_1. For any Lipschitz function $f \in \mathscr{L}$, it is easy to verify f is measurable and the integral $\int f d\mu$ is finite for any $\mu \in \mathscr{P}_{W_1}$.

Let $\phi : \mathscr{P}_{W_1} \to \bar{\mathbb{R}}$. The (Fenchel) conjugate function of ϕ, $\rho : \mathscr{L} \to \bar{\mathbb{R}}$, is defined as

$$\rho(f) := \sup_{\mu \in \mathscr{P}_{W_1}} \left(\int f d\mu - \phi(\mu) \right) \tag{15}$$

and the second conjugate function $\phi^c : \mathscr{P}_{W_1} \to \bar{\mathbb{R}}$ is defined as

$$\phi^c(\mu) := \sup_{f \in \mathscr{L}} \left(\int f d\mu - \rho(f) \right).$$

We can show the following conjugate duality.

THEOREM 4.1. *Let $\phi : \mathscr{P}_{W_1} \to \bar{\mathbb{R}}$ be a convex and lower semicontinuous function on $(\mathscr{P}_{W_1}, W_1)$ satisfying $\phi(\mu) > -\infty$ for all $\mu \in \mathscr{P}_{W_1}$ and $\phi(\mu) \in \mathbb{R}$ for some $\mu \in \mathscr{P}_{W_1}$. Then $\phi(\mu) = \phi^c(\mu)$, $\forall \mu \in \mathscr{P}_{W_1}$.*

PROOF. Since the proof is rather technical, we postpone it to the Appendix B. □

Note that the conjugate function ρ is always convex and has the following properties:

(i) (transition invariance) for any constant function c, $\rho(f + c) = \rho(f) + c$; and

(ii) (monotonicity) for any f, g satisfying $f(x) \leq g(x)$, $\rho(f) \leq \rho(g)$.

This shows that ρ is in fact a *convex risk measure* (or utility functional) which has been widely applied in mathematical finance (see e.g., Föllmer and Schied (2004, Chapter 4)).

Let ϕ be a function on $\mathscr{P}_{W_1}$ and ρ be its conjugate as in (15). Consider the following sets

$$\mathcal{N}_\phi := \{f \in \mathscr{L} | \rho(f) = 0\}, \text{ and}$$
$$\bar{\mathcal{N}}_\phi := \{f \in \mathscr{L} | \rho(f) \leq 0\} = \left\{ f \in \mathscr{L} \middle| \int f d\mu \leq \phi(\mu), \forall \mu \in \mathscr{P}_{W_1} \right\}.$$

We call $\mathcal{N}_\phi$ the *null level-set of* ϕ, whereas the latter set $\bar{\mathcal{N}}_\phi$ is called the *acceptance set of* ϕ (cf. (Föllmer and Schied, 2004, Section 4.1)). Note that since ρ is convex and lower semicontinuous, $\bar{\mathcal{N}}_\phi$ is convex and closed (see e.g. (Ioan-Bot et al., 2009, Theorem 2.2.9)). We have the following dual representation.

COROLLARY 4.2. *Let $\phi : \mathscr{P}_{W_1} \to \bar{\mathbb{R}}$ be a function satisfying the condition in Theorem 4.1. Then*

$$\phi(\mu) = \sup_{f \in \mathcal{N}_\phi} \int f d\mu = \sup_{f \in \bar{\mathcal{N}}_\phi} \int f d\mu.$$

PROOF. (a) We show first $\phi(\mu) = \sup_{f \in \mathcal{N}_\phi} \int f d\mu$. Indeed, Theorem 4.1 yields

$$\phi(\mu) = \sup_{f \in \mathscr{L}} \left(\int f d\mu - \rho(f) \right) \geq \sup_{f \in \mathcal{N}_\phi} \left(\int f d\mu - \rho(f) \right) = \sup_{f \in \mathcal{N}_\phi} \int f d\mu.$$

Since there exists one μ such that $\phi(\mu) \in \mathbb{R}$, we have $\rho(f) > -\infty, \forall f \in \mathscr{L}$. Thus,

$$\sup_{f \in \mathscr{L}} \left(\int f d\mu - \rho(f) \right) = \sup_{f \in \mathscr{L} : \rho(f) < \infty} \left(\int f d\mu - \rho(f) \right) = \sup_{f \in \mathscr{L} : \rho(f) \in \mathbb{R}} \left(\int f d\mu - \rho(f) \right).$$

Due to the translation invariance, $f' := f - \rho(f)$ satisfies that $\rho(f') = 0$ if $\rho(f) \in \mathbb{R}$. Hence,

$$\phi(\mu) = \sup_{f \in \mathscr{L} : \rho(f) \in \mathbb{R}} \left(\int f d\mu - \rho(f) \right) = \sup_{f' : f' = f - \rho(f), f \in \mathscr{L}, \rho(f) \in \mathbb{R}} \int f' d\mu \leq \sup_{f' \in \mathcal{N}_\phi} \int f' d\mu. \tag{16}$$

Combining the above two inequalities yields $\phi(\mu) = \sup_{f \in \mathcal{N}_\phi} \int f d\mu$.

(b) We show now $\phi(\mu) = \sup_{f \in \bar{\mathcal{N}}_\phi} \int f d\mu$. Theorem 4.1 yields

$$\phi(\mu) = \sup_{f \in \mathscr{L}} \left(\int f d\mu - \rho(f) \right) \geq \sup_{f \in \bar{\mathcal{N}}_\phi} \left(\int f d\mu - \rho(f) \right) \geq \sup_{f \in \bar{\mathcal{N}}_\phi} \int f d\mu.$$

On the other hand, (16) yields $\phi(\mu) \leq \sup_{f' \in \mathcal{N}_\phi} \int f' d\mu \leq \sup_{f' \in \bar{\mathcal{N}}_\phi} \int f' d\mu$. □

4.2 Application to POMDPs

LEMMA 4.3. *If $\phi : \mathscr{P}_w \to \mathbb{R}$ is convex, then $\mathcal{T}_a(\phi)$ is convex, $\forall a \in \mathsf{A}$, and therefore, $\mathcal{T}(\phi)$ is convex as well.*

PROOF. It is sufficient to show that $\mu \mapsto r(\mu, a) + \alpha \int \phi(M(\mu, a, y))\tilde{R}(dy|\mu, a)$ is convex for each $a \in \mathsf{A}$. Indeed, take any action $a \in \mathsf{A}$ and let μ_1 and μ_2 be two arbitrary probability measures in $\mathscr{P}_w$. Take any $\kappa \in (0,1)$ and define $\mu_\kappa := \kappa\mu_1 + (1-\kappa)\mu_2$. By the definition, for any $\mathsf{B} \in \mathcal{B}(\mathsf{X})$ and $\mathsf{C} \in \mathcal{B}(\mathsf{Y})$, we have

$$\begin{aligned} R(\mathsf{B}, \mathsf{C}|\mu_\kappa, a) =& \kappa R(\mathsf{B}, \mathsf{C}|\mu_1, a) + (1-\kappa) R(\mathsf{B}, \mathsf{C}|\mu_2, a) && (17)\\ =& \kappa \int_{\mathsf{C}} M(\mathsf{B}|\mu_1, a, y)\tilde{R}(dy|\mu_1, a) + (1-\kappa)\int_{\mathsf{C}} M(\mathsf{B}|\mu_2, a, y)\tilde{R}(dy|\mu_2, a). && (18)\end{aligned}$$

On the other hand, a simple calculation yields

$$\tilde{R}(\mathsf{C}|\mu_\kappa, a) = \kappa\tilde{R}(\mathsf{C}|\mu_1, a) + (1-\kappa)\tilde{R}(\mathsf{C}|\mu_2, a), \forall \mathsf{C} \in \mathcal{B}(\mathsf{Y}). \tag{19}$$

Hence, $\tilde{R}(\mathsf{C}|\mu_\kappa, a) = 0$ implies $\tilde{R}(\mathsf{C}|\mu_1, a) = 0$ and $\tilde{R}(\mathsf{C}|\mu_2, a) = 0, \forall \mathsf{C} \in \mathcal{B}(\mathsf{Y})$. By Radon-Nikodym theorem, there exist functions $f_i : \mathsf{Y} \times \mathsf{A} \to [0, \infty), i = 1, 2$, which are both $\mathcal{B}(\mathsf{Y})$-measurable for the fixed a, such that

$$\tilde{R}(\mathsf{C}|\mu_i, a) = \int_{\mathsf{C}} f_i(y, a)\tilde{R}(dy|\mu_\kappa, a), i = 1, 2. \tag{20}$$

Applying these two equations in (19) accordingly, we obtain

$$\tilde{R}(\mathsf{C}|\mu_\kappa, a) = \int_{\mathsf{C}} \left(\kappa f_1(y, a) + (1-\kappa) f_2(y, a)\right) \tilde{R}(dy|\mu_\kappa, a), \forall \mathsf{C} \in \mathcal{B}(\mathsf{Y}),$$

which implies that $\kappa f_1(y, a) + (1-\kappa) f_2(y, a) = 1$, $\tilde{R}(\cdot|\mu_\kappa, a)$-almost surely. In other words, there exists a Borel set $\bar{\mathsf{C}} \in \mathcal{B}(\mathsf{Y})$ such that $\tilde{R}(\bar{\mathsf{C}}|\mu_\kappa, a) = 1$ and $\kappa f_1(y, a) + (1-\kappa) f_2(y, a) = 1, \forall y \in \bar{\mathsf{C}}$.

Applying (20) to (18), we obtain

$$R(\mathsf{B}, \mathsf{C}|\mu_\kappa, a) = \int_{\mathsf{C}} \left[\kappa M(\mathsf{B}|\mu_1, a, y) f_1(y, a) + (1-\kappa) M(\mathsf{B}|\mu_2, a, y) f_2(y, a)\right] \tilde{R}(dy|\mu_\kappa, a),$$

and for each $\kappa \in (0,1)$, $M(\cdot|a, y, \kappa) := \kappa M(\cdot|\mu_1, a, y) f_1(y, a) + (1-\kappa) M(\cdot|\mu_2, a, y) f_2(y, a)$ is a valid stochastic kernel satisfying $R(\mathsf{B}, \mathsf{C}|\mu_\kappa, a) = \int_{\mathsf{C}} M(\mathsf{B}|a, y, \kappa)\tilde{R}(dy|\mu_\kappa, a), \forall \mathsf{B} \in \mathcal{B}(\mathsf{X}), \mathsf{C} \in \mathcal{B}(\mathsf{Y})$. Finally, the convexity of ϕ implies

$$\begin{aligned} &\int \phi(M(\cdot|a, y, \kappa))\tilde{R}(dy|\mu_\kappa, a) = \int_{\bar{\mathsf{C}}} \phi(M(\cdot|a, y, \kappa))\tilde{R}(dy|\mu_\kappa, a)\\ \leq& \int_{\bar{\mathsf{C}}} \left[\kappa f_1(y, a)\phi(M(\mu_1, a, y)) + (1-\kappa) f_2(y, a)\phi(M(\mu_2, a, y))\right] \tilde{R}(dy|\mu_\kappa, a)\\ \leq& \kappa \int \phi(M(\mu_1, a, y))\tilde{R}(dy|\mu_1, a) + (1-\kappa)\int \phi(M(\mu_2, a, y))\tilde{R}(dy|\mu_2, a),\end{aligned}$$

which yields the required convexity. □

We introduce the following assumption accompanying Assumption 3.5.

Assumption 4.4. For each $a \in \mathsf{A}$, there exist stochastic kernels M on $\mathscr{P}$ given $\mathscr{P}_w \times \mathsf{A} \times \mathsf{Y}$ and $\tilde{R}$ on Y given $\mathscr{P}_w \times \mathsf{A}$ satisfying (7) such that, if $\{\mu_n \in \mathscr{P}_w, n = 1, 2, \ldots\}$ converges to $\mu_0 \in \mathscr{P}_w$ as $n \to \infty$,

(i) there exists a subsequence $\{\mu_{n_k}\} \subset \{\mu_n\}$ and a measurable set $\bar{\mathsf{C}} \in \mathcal{B}(\mathsf{Y})$ such that $\tilde{R}(\bar{\mathsf{C}}|\mu_0, a) = 1$ and for all $y \in \bar{\mathsf{C}}$, $M(\mu_{n_k}, a, y)$ converges canonically weakly to $M(\mu_0, a, y)$;

(ii) for each $\mathsf{C} \in \mathcal{B}(\mathsf{Y})$, $\tilde{R}(\mathsf{C}|\mu_n, a) \to \tilde{R}(\mathsf{C}|\mu_0, a)$ as $n \to \infty$.

LEMMA 4.5. *Suppose Assumptions 4.4, 3.1(i) and 3.3(i) hold. Then, for any $\phi \in \mathscr{B}_{\tilde{w}}$, $\mathcal{T}_a(\phi)$ is continuous for any $a \in \mathsf{A}$ and therefore $\mathcal{T}(\phi)$ is lower semicontinuous.*

PROOF. Fix one $a \in \mathsf{A}$. Note that Proposition 3.2(iii) ensures the continuity of the function $\mu \mapsto r(\mu, a)$. It remains to show that $\mu \mapsto \int \phi(\mu')\tilde{Q}(d\mu'|\mu, a)$ is continuous.

By Assumption 4.4 and (Feinberg et al., 2016, Theorem 3.4), $\mu \mapsto \tilde{Q}(d\mu'|\mu, a)$ is canonically weakly continuous. By Assumption, there exists a constant $\bar{\phi} > 0$ such that $|\phi(\mu)| \leq \bar{\phi}\tilde{w}(\mu) = \bar{\phi}\int w d\mu$. Let $\phi'(\mu) := \phi(\mu) + \bar{\phi}\int w d\mu$, which is nonnegative. Hence, it is a limit of a nondecreasing sequence of measurable bounded function $\{\phi'_m\}$ such that $\phi'_m \uparrow \phi'$. Let $\{\mu_n \in \mathscr{P}_w\}$ be a converging sequence with under the W_1-metric to a limit $\mu_0 \in \mathscr{P}_w$. We have then

$$\liminf_{n\to\infty} \int \phi'(\mu')\tilde{Q}(d\mu'|\mu_n, a) \geq \liminf_{n\to\infty} \int \phi'_m(\mu')\tilde{Q}(d\mu'|\mu_n, a) = \int \phi'_m(\mu')\tilde{Q}(d\mu'|\mu_0, a).$$

Hence, letting $m \to \infty$, monotone convergence yields that

$$\liminf_{n\to\infty} \int \phi'(\mu')\tilde{Q}(d\mu'|\mu_n, a) \geq \int \phi'(\mu')\tilde{Q}(d\mu'|\mu_0, a). \tag{21}$$

On the other hand, we have for each $n \in \mathbb{N}$,

$$\int \tilde{w}(\mu')\tilde{Q}(d\mu'|\mu_n, a) = \int_{\mathsf{Y}}\int_{\mathsf{X}} w(x')\tilde{P}(dx'|\mu_n, a)Q(dy|x', a) = \int_{\mathsf{X}}\int_{\mathsf{X}} w(x')P(dx'|x, a)\mu_n(dx).$$

Note that by Assumption 3.3(i), we have for each $a \in \mathsf{A}$, $w'(x, a) := \int_{\mathsf{X}} w(x')P(dx'|x, a) \in \mathscr{L}_w$. Proposition 2.4(ii) yields that $\lim_{n\to\infty}\int \tilde{w}(\mu')\tilde{Q}(d\mu'|\mu_n, a) = \int \tilde{w}(\mu')\tilde{Q}(d\mu'|\mu_0, a)$. Hence, (21) implies that $\liminf_{n\to\infty}\int \phi(\mu')\tilde{Q}(d\mu'|\mu_n, a) \geq \int \phi(\mu')\tilde{Q}(d\mu'|\mu_0, a)$. In other words, $\mu \mapsto \int \phi(\mu')\tilde{Q}(d\mu'|\mu, a)$ is lower semicontinuous. We apply this fact to $-\phi$ in lieu of ϕ and obtain that $\mu \mapsto \int \phi(\mu')\tilde{Q}(d\mu'|\mu, a)$ is also upper semicontinuous. Thus the required continuity holds. □

We immediately obtain the following result.

THEOREM 4.6. *Suppose Assumptions 4.4, 3.1(i) and 3.3(i) hold. If $\phi \in \mathscr{B}_{\tilde{w}}$ is convex, then $\mathcal{T}(\phi)$ is convex and lower semicontinuous.*

Remark 4.7. Feinberg et al. (2016, Theorem 3.6) show that one sufficient condition to guarantee both Assumption 3.5 and 4.4 is that (i) The stochastic kernel $P(dx'|x, a)$ is canonically weakly continuous and (ii) the stochastic kernel $Q(dy|x, a)$ is continuous in total variation. In addition, it is demonstrated in (Feinberg et al., 2016, Example 4.1) that the latter continuity in total variation cannot be weakened to the canonical weak continuity.

4.3 Set iteration

Recall that Corollary 4.2 imply that a convex and lower semicontinuous function ϕ admits a representation of $\phi(\mu) = \sup_{f\in\mathcal{N}} \inf \int f d\mu$ with some set $\mathcal{N} \subset \mathscr{L}$. Hence, instead of iterating the value function, we can iterate the acceptance set, which is described as follows.

Algorithm 4.8. Start with any set $\bar{\mathcal{N}}_0 \subset \mathscr{L}$. At time t, update the acceptance set using the following two steps:

$$\phi_{t+1}(\mu) = \max_{a\in\mathsf{A}} \left(r(\mu, a) + \alpha \int \left(\sup_{f\in\bar{\mathcal{N}}_t} \int f(x')M(dx'|\mu, a, y) \right) \tilde{R}(dy|\mu, a) \right)$$

$$\bar{\mathcal{N}}_{t+1} = \left\{ f \in \mathscr{L} \;\middle|\; \phi_{t+1}(\mu) \geq \int f d\mu, \forall \mu \in \mathscr{P}_w \right\}.$$

These steps will be repeated until some stopping criterion is satisfied.

An iteration of null level-sets can be analogously designed as above and is therefore omitted. Note that in course of iteration, Lemma 4.3 and 4.5 guarantee that ϕ_t is convex and lower semicontinuous for each $t = 1, 2, \ldots$. Hence, Corollary 4.2 ensures that $\phi_t(\mu) = \sup_{f\in\bar{\mathcal{N}}_t} \inf f d\mu, \forall \mu \in \mathscr{P}_w$, and for each $t = 1, 2, \ldots$. By Theorem 3.8(a), we immediately obtain the following result.

THEOREM 4.9. *Suppose Assumptions 3.1, 3.3, 3.5, 3.7 and 4.4 hold. Let ϕ^* be the optimal value function for the POMDP, $\bar{r} > 0$ and $\gamma \in (0,1)$ be the constants as in Theorem 3.8(a). Then,*

$$\|\phi^* - \phi_t\|_{\tilde{w}} \leq \bar{r}\gamma^t/(1-\gamma), \text{ where } \phi_t(\mu) = \sup_{f \in \bar{\mathcal{N}}_t} \int f d\mu, \forall \mu \in \mathscr{P}_w.$$

This implies that the optimal value function ϕ^* can be arbitrarily well approximated by a convex and lower semicontinuous function ϕ of the dual form.

COROLLARY 4.10. *Suppose Assumptions 3.1, 3.3, 3.5, 3.7 and 4.4 hold. For any $\epsilon > 0$, there exists a set $\mathcal{N}^\epsilon \subset \mathscr{L}$ satisfying $\|\phi^* - \phi^\epsilon\|_{\tilde{w}} \leq \epsilon$, where $\phi^\epsilon(\mu) := \sup_{f \in \mathcal{N}^\epsilon} \int f d\mu, \forall \mu \in \mathscr{P}_w$.*

Similar to Q-value iteration in MDP literature (see e.g. Hernández-Lerma and Lasserre (1996)), we can iterate acceptance sets depending on a.

Algorithm 4.11. Start with an initial set $\bar{\mathcal{N}}_0 \subset \mathscr{L}$ and set $\bar{\mathcal{N}}_0^a = \bar{\mathcal{N}}_0, \forall a \in \mathsf{A}$. At time t, update the null level set using the following two steps:

$$\phi_{t+1}(\mu, a) = \tilde{r}(\mu, a) + \alpha \int \left(\sup_{f \in \bigcup_{a \in \mathsf{A}} \bar{\mathcal{N}}_t^a} \int f(x') M(dx'|\mu, a, y) \right) \tilde{R}(dy|\mu, a)$$

$$\bar{\mathcal{N}}_{t+1}^a = \left\{ f \in \mathscr{L} \;\middle|\; \phi_{t+1}(\mu, a) \geq \int f d\mu, \forall \mu \in \mathscr{P}_w \right\}.$$

These steps will be repeated until some stopping criterion is satisfied.

Under Assumptions 3.1, 3.3, 3.5, 3.7, the iteration stated in Algorithm 4.11 satisfies

$$\|\phi^*(\cdot) - \max_{a \in \mathsf{A}} \phi_t(\cdot, a)\|_{\tilde{w}} \leq \bar{r}\gamma^t/(1-\gamma).$$

Q with a reference measure Let us assume that there exists a reference (probability) measure φ on Y such that $Q(\cdot|x', a) \ll \varphi(\cdot)$ for all $(x', a) \in \mathsf{X} \times \mathsf{A}$. Note that POMDPs in many applications satisfy this assumption. For example, the assumption holds automatically if the observation space is finite. Hence, the density of Q w.r.t. φ exists and is denoted by $q(y|x', a)$. In this case, the iteration can be further simplified. Indeed, it is easy to verify that

$$M(dx'|\mu, a, y) = \frac{\tilde{P}(dx'|\mu, a) q(y|x', a)}{\int_{\mathsf{X}} \tilde{P}(dx'|\mu, a) q(y|x', a)} \quad \text{and} \quad \tilde{R}(dy|\mu, a) = \int_{\mathsf{X}} \tilde{P}(dx'|\mu, a) q(y|x', a) \varphi(dy)$$

satisfy (7). Under this setup, the calculation of iteration becomes much simpler. Suppose $\phi \in \mathscr{B}_{\tilde{w}}$ is convex and lower semicontinuous, then we have by Corollary 4.2,

$$\mathcal{T}_a(\phi)(\mu) = r(\mu, a) + \alpha \int \left(\sup_{f \in \bar{\mathcal{N}}_\phi} \int f(x') \tilde{P}(dx'|\mu, a) q(y|x', a) \right) \varphi(dy). \tag{22}$$

In particular, the continuity of Q in total variation mentioned in Remark 4.7 is

$$\int |q(y|x_n, a_n) - q(y|x_0, a_0)| \varphi(dy) \to 0, \text{ as } (x_n, a_n) \to (x_0, a_0).$$

5 One example

In this section, we apply our results to the nonlinear Kalman filter in discrete time, which has been widely used in signal processing and finance (see e.g. Bain et al. (2009)). To adapt it to the POMDP-framework, we introduce a control variable in the dynamics of both hidden states and observations. We consider the one-dimensional case only, but results can be easily extended to multidimensional cases. By this example we would like to demonstrate, that the assumptions made in the previous sections can be satisfied for a properly chosen metric.

Consider a 1-dimensional simple linear model

$$x_{t+1} = d + b(a_t) x_t + \sigma n_t$$

where n_t is i.i.d. white noise with $\sigma > 0$ being the standard deviation, d is a real constant, and $b : \mathsf{A} \to \mathbb{R}$ is a continuous function on A satisfying $\sup_{a\in\mathsf{A}}|b(a)| = \epsilon < 1$. Hence, the transition kernel P is given by

$$P(dx'|x,a) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x'-d-b(a)x)^2}{2\sigma^2}} dx'.$$

We assume that the observation dynamic is given by

$$y_t = h(x_t, a_t) + \tilde{\sigma}\tilde{n}_t$$

where $h : \mathbb{R} \times \mathsf{A} \to \mathbb{R}$ is a continuous function on $\mathbb{R} \times \mathsf{A}$, $\tilde{n}_t$ is i.i.d. white noise which is independent of n_t with $\tilde{\sigma} > 0$ being the standard deviation. Hence, Q is given by

$$Q(dy|x,a) = \frac{1}{\sqrt{2\pi}\tilde{\sigma}} e^{-\frac{(y-h(x,a))^2}{2\tilde{\sigma}^2}} dy.$$

Let $d(x, x_0) = |x - x_0|$ be the Euclidean distance on $\mathbb{R}$ and set $x_0 = 0$. Then

$$\left(\int d(x', x_0)P(dx'|x,a)\right)^2 \le \int d^2(x', x_0)P(dx'|x,a) = \sigma^2 + (d + b(a)x)^2$$

Since $b^2(a) \le \epsilon^2 < 1$ for all $a \in \mathsf{A}$, there exists a $\tilde{\beta} \in (\epsilon, 1)$ and $K > 0$ such that

$$\int d(x', x_0)P(dx'|x,a) \le \tilde{\beta}d(x, x_0) + K, \forall x \in \mathbb{R}.$$

For any discount factor $\alpha \in (0, 1)$, we may select any $\beta \in (1, 1/\alpha)$ and set $k := \frac{\beta-1}{K}$. The above inequality yields

$$\begin{aligned}\int \left(1 + kd(x', x_0)\right) P(dx'|x,a) \le& 1 + \frac{\beta-1}{K}(\tilde{\beta}d(x, x_0) + K)\\ \le& \tilde{\beta}(1 + kd(x, x_0)) \le \beta(1 + kd(x, x_0)).\end{aligned}$$

Hence, the Assumption 3.3(i) holds. It is easy to verify that Assumption 3.3(ii) holds due to the continuity of b.

For the reward function, we assume that there exists a positive constant $\tilde{r} < \infty$ such that $|r(x,a)| \le \tilde{r}(1 + d(x, x_0)) = \tilde{r}(1 + |x|)$ holds for all $(x,a) \in \mathsf{K}$. Furthermore, we assume that $a \mapsto r(x,a)$ is upper semicontinuous. Then, it is easy to verify Assumption 3.1. It is worth to be mentioned that the reward function in this example is allowed to be unbounded on both sides.

Finally, it remains to verify Assumption 3.5 and 4.4. By Remark 4.7 (see Feinberg et al. (2016, Theorem 3.6)), it is sufficient to show that $P(\cdot|x,a)$ is canonically weakly continuous and $Q(dy|x,a)$ is continuous in total variation. The weak continuity is easy to verify due to the continuity of b. We only show the continuity of $Q(dy|x,a)$ in total variation. Let $q(y,x) := \frac{1}{\sqrt{2\pi}\tilde{\sigma}} e^{-\frac{(y-h(x))^2}{2\tilde{\sigma}^2}}$. Hence, it is to show that

$$\int_{-\infty}^{\infty} |q(y, x_n) - q(y, x_0)|\, dy = \int_{-\infty}^{\infty} \left|\frac{q(y, x_n)}{q(y, x_0)} - 1\right| q(y, x_0)dy \to 0, \text{ as } n \to \infty. \tag{23}$$

Indeed, since h is continuous, we have for each $y \in \mathbb{R}$, $\lim_{n\to\infty} \frac{q(y,x_n)}{q(y,x_0)} = 1$. On the other hand, we have

$$\frac{q(y, x_n)}{q(y, x_0)} \le e^{\frac{h^2(x_0)-h^2(x_n)}{2\tilde{\sigma}^2}} e^{\frac{\epsilon^2 y^2}{2\tilde{\sigma}^2} + \frac{(h(x_0)-h(x_n))^2}{2\epsilon^2\tilde{\sigma}^2}}, \forall \epsilon > 0,$$

and $|h(x_n)|$ is uniformly bounded w.r.t. n due to the continuity of h. Since ϵ can be taken any value, we may select some ϵ such that $\int e^{\frac{\epsilon^2 y^2}{2\tilde{\sigma}^2}} q(y, x_0) < \infty$. Hence, there exists a function $g(y)$ such that $\frac{q(y,x_n)}{q(y,x_0)} \le g(y)$ and $\int g(y)q(y, x_0)dy < \infty$. By Fatou-Lebesgue theorem, we obtain the required convergence stated in (23).

A Markov decision processes

For readers' convenience, we state below an adapted version of the classical results of Markov decision processes (MDPs) on weighted norm spaces. For the original statements, we refer to Hernández-Lerma and Lasserre (1999, Chapter 8).

Let $\mathcal{M} := \{\mathsf{X}, \mathsf{A}, Q, r\}$ be an MDP, where X and A are Borel spaces denoting the state and action space respectively. $Q(\cdot|x,a)$ denotes the transition kernel, which is a stochastic kernel from $\mathsf{X}\times\mathsf{A}$ to X, where $\mathsf{X}\times\mathsf{A}$ is also a Borel space. $r : \mathsf{X}\times\mathsf{A}$ is the reward function, which is $\mathcal{B}(\mathsf{X}\times\mathsf{A})$-measurable.

For each $t = 0, 1, 2, \ldots$, let H_t be the family of history up to time t, i.e., $\mathsf{H}_0 = \mathsf{X}$ and $\mathsf{H}_t = (\mathsf{X}\times\mathsf{A})^t \times \mathsf{X} = \mathsf{X}\times\mathsf{A}\times\mathsf{H}_{t-1}$ for $t = 1, 2, \ldots$. An element of H_t is a vector of the form $h_t = (x_0, a_0, \ldots, x_{t-1}, a_{t-1}, x_t)$. Let $\boldsymbol{\pi} := [\pi_0, \pi_1, \ldots]$ be a deterministic policy where $\pi_t : \mathsf{H}_t \to \mathsf{A}$ denotes the policy at time t and Π be the space of all deterministic policies. A policy is called Markov, if there exists a sequence $\{\varphi_t : \mathsf{X}\to\mathsf{A}\}$ of measurable functions such that $\pi_t(h_t) = \varphi_t(x_t), \forall h_t \in \mathsf{H}_t, t = 0, 1, 2, \ldots$. It is well known that the policy set can be extended to include also randomized policies, but the following results remain the same. Then, the expected discount reward starting from any $x \in \mathsf{X}$ is

$$J(x, \boldsymbol{\pi}) := \mathbb{E}^{\boldsymbol{\pi}}\left[\sum_{t=0}^{\infty} \alpha^t r(x_t, a_t)\middle| x_0 = x\right]$$

and its optimal value function is

$$v^*(x) := \sup_{\boldsymbol{\pi}\in\Pi} J(x, \boldsymbol{\pi}), \forall x \in \mathsf{X}.$$

Let $v : \mathsf{X} \to \mathbb{R}$ be a function and consider the following operator:

$$\mathcal{T}v(x) := \sup_{a\in\mathsf{A}}\left(r(x,a) + \alpha\int v(x')Q(dx'|x,a)\right)$$

Note that under proper assumptions, the optimal a^* in the above optimization problem is attainable in A for each $x \in \mathsf{X}$ and therefore, we can replace "sup" by "max" in the above definition. Then, the value iteration works as follows. We start from arbitrary v_0. At time t, the value function is updated as $v_{t+1} = \mathcal{T}(v_t)$

Assumption A.1. (a) The action set A is compact; For each $x \in \mathsf{X}$, (b) the reward function $r(x,a)$ is upper semicontinuous in $a \in \mathsf{A}$; (c) for each bounded continuous function $f : \mathsf{X}\to\mathbb{R}$, $a \mapsto \int f(y)Q(dy|x,a)$ is continuous.

Assumption A.2. There exist a weight function $w : \mathsf{X}\to[1,\infty)$ and a constant $\beta \in (0, 1/\alpha)$ such that (a) $\sup_{a\in\mathsf{A}}|r(x,a)| \le \bar{r}w(x), \forall x \in \mathsf{X}$, with some constant $\bar{r} > 0$; (b) $\sup_{a\in\mathsf{A}}\int w(y)Q(dy|x,a) \le \beta w(x), \forall x \in \mathsf{X}$.

Assumption A.3. For each $x \in \mathsf{X}$, $a \mapsto \int w(y)Q(dy|x,a)$ is continuous.

We restate the main result from MDPs (see Hernández-Lerma and Lasserre (1999, Theorem 8.3.6)) as follows

Theorem A.4. *Suppose that Assumptions A.1 – A.3 hold. Let β be the constant in Assumption A.2(b) and $\bar{r}$ be the constant in A.2(a) and define $\gamma := \alpha\beta \in (0,1)$. Then*

(a) the optimal value function $v^* \in \mathscr{B}_w$ is the unique fixed point of the operator $\mathcal{T}$ satisfying

$$v^*(x) = \mathcal{T}(v^*)(x) = \max_{a\in\mathsf{A}}\left(r(x,a) + \alpha\int \phi\left(x'\right) Q(dx'|x,a)\right)$$

and $\|v_t - v^*\|_w \le \bar{r}\gamma^t/(1-\gamma), t = 1, 2, \ldots$.

(b) there exists a selector $f^* : \mathsf{X}\to\mathsf{A}$ such that

$$v^*(x) = r(x, f^*(x)) + \alpha\int \phi\left(x'\right) Q(dx'|x, f^*(x)), \forall x \in \mathsf{X},$$

and $\boldsymbol{\pi}^* = (f^*)^\infty$ is an optimal policy that is Markov and satisfies $v^*(x) = J(x, \boldsymbol{\pi}^*)$ for each $x \in \mathsf{X}$.

B Dual representation of functions on the Wasserstein space

Dual representation, supported by the Fenchel-Moreau theorem (see e.g., Ioan-Bot et al. (2009); Zălinescu (2002)), plays an important role in convex analysis and has wide applications in various fields (see e.g., Borwein and Lewis (2006); Boyd and Vandenberghe (2004)). In this paper, we are interested in the conjugate dual representation of functions on probability measures. Unfortunately, we cannot apply Fenchel-Moreau theorem in this case, because it is difficult to identify the dual space of all probability measures, $\mathscr{P}$, endowed with the usual topology, e.g., the weak topology. In the following, we consider the topology induced the Wasserstein metric (see e.g., (Villani, 2009, Chapter 6)), which yields nice topological properties. Under this setting, we shall prove a dual representation of the Fenchel-Moreau type.

B.1 Wasserstein space

Let (X, d) be a Polish space. Let W_1 be the Wasserstein metric defined in (11) and $\mathscr{P}_{W_1}$ be the Wasserstein space of order 1. We state below the key property of Wasserstein spaces.

THEOREM B.1. *Let (X, d) be a Polish space. Then, $(\mathscr{P}_{W_1}, W_1)$ is also a Polish space, i.e., a complete separable metric space. In particular, let $D := \{x_i, i = 1, 2, \ldots\}$ be a dense subset of X, then for any $\mu \in \mathscr{P}_{W_1}$ and $\epsilon > 0$, there exists a subset of $\{x_k, k = 1, 2, \ldots, n\} \subset D$ and rational numbers $\{a_k, k = 1, 2, \ldots, a_k \geq 0\}$ such that $W_1(\mu, \sum_{k=1}^n a_k \delta_{x_k}) < \epsilon$, where $\{a_k\}$ satisfying $\sum_{k=1}^n a_k = 1$ and δ_x denotes the Dirac measure at x.*

PROOF. See the proof of Theorem 6.18 in Villani (2009). □

We first consider the dense subset of $\mathscr{P}_{W_1}$. Let $D := \{x_i, i = 1, 2, \ldots\}$ be a dense subset of X. Let $\mathscr{D}$ be the space of all probability measures in $\mathscr{P}$ with a finite support, which is contained in D, i.e.,

$$\mathscr{D} := \left\{ \sum_{i=1}^n a_i \delta_{x_i} \middle| n \in \mathbb{N}, a_i \in \mathbb{R}_+ i = 1, 2, \ldots, n, \sum_{i=1}^n a_i = 1 \right\}, \tag{24}$$

By Theorem B.1, $\mathscr{D}$ is dense in $\mathscr{P}_{W_1}$ equipped with the Wasserstein metric W_1. Replacing the Dirac measure with the indicator function

$$\mathbf{1}_x(y) := \begin{cases} 1, & x = y \\ 0, & x \neq y \end{cases}$$

in $\mathscr{D}$, we define the following set of real-valued functions on X

$$\mathcal{D} := \left\{ \sum_{i=1}^n a_i \mathbf{1}_{x_i} \middle| n \in \mathbb{N}, a_i \in \mathbb{R}_+, i = 1, 2, \ldots, n, \sum_{i=1}^n a_i = 1 \right\}. \tag{25}$$

Apparently $\mathscr{D}$ and $\mathcal{D}$ are isometric and let $\psi : \mathscr{D} \to \mathcal{D}$ be the bijective mapping. Fix a point $\nu_0 \in \mathscr{D}$. Then we have

$$\sum_{x \in D} \psi(\nu)(x) - \psi(\nu_0)(x) = 0, \forall \nu \in \mathscr{D}.$$

We will see below that $\psi(\nu) - \psi(\nu_0)$ is in fact an element of the *Arens-Eells space*, whose dual space is exactly the space of all Lipschitz functions, $\mathscr{L}$.

B.2 Arens-Eells space

We recall some results of the Arens-Eells space based on Weaver (1999, Section 2.2 and 2.3).

DEFINITION B.2. *Let (X, d) be a metric space. A* molecule *of X is a function $m : \mathsf{X} \to \mathbb{R}$ which is supported on a finite set and which satisfies $\sum_{x \in \mathsf{X}} m(x) = 0$.*

For any $x, y \in \mathsf{X}$ define the molecule $m_{xy} = \mathbf{1}_x - \mathbf{1}_y$, where $\mathbf{1}_x$ denotes the indicator function on the singleton set $\{x\}$. Define the following seminorm

$$\|m\|_{Æ} := \inf \left\{ \sum_{i=1}^{n} |a_i| d(x_i, y_i) \middle| m = \sum_{i=1}^{n} a_i m_{x_i y_i} \right\}$$

and let Æ be the completion of the space of molecules, which is also called *Arens-Eells sapce*. The following theorem states that Æ is one predual of $\mathscr{L}$, the space of Lipschitz functions.

THEOREM B.3 (Theorem 2.2.2 Weaver (1999)). *Let (X, d) be a metric space with at least one point x_0. Then the space of Lipschitz functions $\mathscr{L}$ is isometrically isomorphic to Æ*, the dual space of Æ, i.e., Æ* $= \mathscr{L}$. In particular, for any $f \in$ Æ*, $\|f\| = \|f\|_{\mathrm{Lip}}$.*

For any $f \in \mathscr{L}$ and $m \in$ Æ, we define

$$\langle f, m \rangle := \sum_{x \in \mathsf{X}} f(x) m(x).$$

COROLLARY B.4. *(i) $\|m\|_{Æ} = \max_{f \in [\mathscr{L}]_1} \langle f, m \rangle$. (ii) $\|\cdot\|_{Æ}$ is a norm on Æ.*

PROOF. See (Weaver, 1999, Corollary 2.2.3). □

Recall that $\mathscr{D}$ defined in (24) is the space of all probability measures in $\mathscr{P}$ with a finite support, which is contained in a dense subset $D \subset \mathsf{X}$, and $\mathcal{D}$ defined in (25) is its corresponding space of functions on X, with $\psi : \mathscr{D} \to \mathcal{D}$ being the bijective map.

PROPOSITION B.5. *$W_1(\nu, \nu_0) = \|\psi(\nu) - \psi(\nu_0)\|_{Æ}$, for any $\nu \in \mathscr{D}$ and $\nu_0 \in \mathscr{D}$.*

PROOF. By the dual representation (12) of W_1, we have

$$\begin{aligned} W_1(\nu, \nu_0) &= \sup_{f \in [\mathscr{L}]_1} \int f d\nu - \int f \nu_0 \\ &= \sup_{f \in [\mathscr{L}]_1} \sum_{x \in \mathsf{X}} \left(\psi(\nu)(x) - \psi(\nu_0)(x) \right) f(x) = \|\psi(\nu) - \psi(\nu_0)\|_{Æ}, \end{aligned}$$

where the last equality is due to Corollary B.4(i). □

B.3 A separation theorem on $\mathscr{D}$

We are now ready to state a separation theorem on $\mathscr{D}$. We first restate the Nirenberg-Luenberger theorem (see Luenberger (1969, Section 5.13), also known as the *minimum norm duality theorem*).

DEFINITION B.6. *Let K be a convex set in a real normed vector space X. Let X^* be its dual space. The function $h(x^*) = \sup_{x \in K} \langle x, x^* \rangle$ on X^* is called the support functional of K.*

THEOREM B.7. *Let x_1 be a point in a real normed space X and let $d > 0$ denotes its distance from the convex set K having support functional h. Then*

$$d = \inf_{x \in K} \|x - x_1\| = \max_{\|x^*\| \leq 1} \left[\langle x_1, x^* \rangle - h(x^*) \right]$$

where the maximum on the right is achieved by some $x_0^* \in X^*$.

THEOREM B.8. *Let K be a convex set of $\mathscr{D}$ and $\nu_0 \in \mathscr{D}$ be a point not contained in K satisfying $W_1(\nu, \nu_0) \geq \epsilon > 0, \forall \nu \in K$. Then there exists a function $f \in [\mathscr{L}]_1$ such that $\int f d\nu \geq \int f d\nu_0 + \epsilon, \forall \nu \in K$.*

PROOF. Since K is convex, $K' := \{\psi(\nu) - \psi(\nu_0) \mid \nu \in K\}$ is also convex in Æ. To apply Theorem B.7, we set $X =$ Æ and its dual is $X^* =$ Æ* $= \mathscr{L}$. By the definition of h, we have

$$h(f) = \sup_{m \in K'} \langle f, m \rangle = \sup_{\nu \in K} \left\{ \int f d\nu - \int f d\nu_0 \right\}, f \in \mathscr{L}.$$

Then by Theorem B.7, we obtain

$$\epsilon \leq \inf_{\nu \in K} W_1(\nu, \nu_0) = \inf_{m \in K'} \|m\|_{Æ} = \max_{f \in [\mathscr{L}]_1} \left[-h(f) \right],$$

which yields $\max_{f\in[\mathscr{L}]_1}\left[\inf_{\nu\in K}\int(-f)d\nu-\int(-f)d\nu_0\right]\geq\epsilon$. Suppose the maximum is attained at f_0. Then $\inf_{\nu\in K}\int(-f_0)d\nu-\int(-f_0)d\nu_0\geq\epsilon$. which implies that $\int(-f_0)d\nu\geq\int(-f_0)d\nu_0+\epsilon,\forall\nu\in K$. □

In later sections, we need to consider the dual space of the Cartesian product $Æ\times\mathbb{R}$, where we can obtain a similar result as above. The operator $\vee$ is defined as $a\vee b:=\max(a,b)$.

THEOREM B.9. *Let $\tilde{K}$ be a convex set of $\mathscr{D}\times\mathbb{R}$ and $(\nu_0,r_0)\in\mathscr{D}\times\mathbb{R}$ be a point not contained in K. Suppose $W_1(\nu,\nu_0)\vee|r-r_0|\geq\epsilon>0,\forall(\nu,r)\in\tilde{K}$. Then there exists a function $(f,\alpha)\in[\mathscr{L}]_1\times[-1,1]$ satisfying $\|f\|_{\text{Lip}}+|\alpha|\leq 1$ such that $\int fd\nu+\alpha r\geq\int fd\nu_0+\alpha r_0+\epsilon,\forall(\nu,r)\in\tilde{K}$.*

PROOF. Let $Æ\times\mathbb{R}$ be equipped with the canonical norm $\|(m,r)\|:=\|m\|_{Æ}\vee|r|$. An easy extension of Theorem B.3 shows that its dual space is isometric with $\mathscr{L}\times\mathbb{R}$ with the norm $\|(f,r)\|=\|f\|_{\text{Lip}}+|r|$, i.e., $(Æ\times\mathbb{R})^*=\mathscr{L}\times\mathbb{R}$. The rest of the proof follows similarly as in the proof of Theorem B.8 and is therefore omitted. □

B.4 A separation theorem on Wasserstein space

We extend the separation Theorem B.8 to the whole Wasserstein space $\mathscr{P}_{W_1}$.

THEOREM B.10. *Let A be a convex and closed subset in $\mathscr{P}_{W_1}$ under the Wasserstein metric W_1. Let $\mu_0\in\mathscr{P}_{W_1}$ be a point not contained in A. Then there exists a function $f\in\mathscr{L}$ such that $\int fd\mu>\int fd\mu_0,\forall\mu\in A$.*

PROOF. Step 1. We first show that there exists a positive constant $\epsilon_0>0$ satisfying $W_1(\mu,\mu_0)\geq\epsilon_0,\forall\mu\in A$. Indeed, suppose the opposite: for any $\epsilon>0$, there exists some $\mu\in A$ satisfying $W_1(\mu,\mu_0)<\epsilon$. Then we can find a set of $\{\mu_n\in A,n=1,2,\ldots\}$ satisfying $W_1(\mu_n,\mu_0)<1/n,n=1,2,\ldots$. Hence, $\mu_n\to\mu_0$ implies that $\mu_0\in A$ due to the fact that A is closed. It contradicts the assumption $\mu_0\notin A$.

Step 2. Let $D:=\{x_1,x_2,\ldots\}$ be a dense set in X and $\mathscr{D}$ as defined in (24), i.e., the space of all probability measures with finite support contained in D. By Theorem B.1, we can approximate μ by a point $\nu\in\mathscr{D}$ with any accuracy. Furthermore, $\mathscr{D}$ is essentially embedded in Æ, whose dual space is known to be $\mathscr{L}$. Finally, Theorem B.8 can be applied to find the required separation function in $\mathscr{L}$. More specifically, take $\epsilon_1=\epsilon_0/5$ and define subsets of $\mathscr{D}$ as follows

$$B_{\epsilon_1}(\mu):=\{\nu\in\mathscr{D}\mid W_1(\mu,\nu)<\epsilon_1\},A_{\epsilon_1}:=\bigcup_{\mu\in A}B_{\epsilon_1}(\mu).$$

Since $\mathscr{D}$ is dense in $(\mathscr{P}_{W_1},W_1)$, $B_{\epsilon_1}(\mu)\neq\emptyset$, $\forall\mu\in\mathscr{P}_{W_1}$. Hence, for any $\nu_i\in A_{\epsilon_1}$ there exists a $\mu_i\in\mathscr{P}_{W_1}$ such that $W_1(\nu_i,\mu_i)<\epsilon_1$, $i=1,2$. Let $\nu_\alpha:=\alpha\nu_1+(1-\alpha)\nu_2$ and $\mu_\alpha:=\alpha\mu_1+(1-\alpha)\mu_2$, $\alpha\in(0,1)$. We have, by the dual representation (12) of W_1,

$$W_1(\nu_\alpha,\mu_\alpha)\leq\alpha W_1(\nu_1,\mu_1)+(1-\alpha)W_1(\nu_2,\mu_2)<\epsilon_1,$$

which implies that A_{ϵ_1} is convex. Similarly, we can find a $\nu_0\in\mathscr{D}$ such that $W_1(\mu_0,\nu_0)<\epsilon_1$. Note that $W_1(\nu,\mu)\geq W_1(\mu_0,\mu)-W_1(\mu_0,\nu_0)-W_1(\nu_0,\nu)\geq\epsilon_0-\epsilon_1-W_1(\nu_0,\nu),\forall\nu\in A_{\epsilon_1},\mu\in A$, yields $\inf_{\mu\in A}W_1(\nu,\mu)\geq\epsilon_0-\epsilon_1-W_1(\nu_0,\nu),\forall\nu\in A_{\epsilon_1}$. By $\epsilon_1\geq\inf_{\mu\in A}W_1(\nu,\mu)$, we have $W_1(\nu_0,\nu)\geq\epsilon_0-2\epsilon_1>0,\forall\nu\in A_{\epsilon_1}$. Hence, by Theorem B.8, there exists a function $f_0\in\mathscr{L}:\|f_0\|_{\text{Lip}}\leq 1$ such that

$$\int f_0d\nu\geq\int f_0d\nu_0+\epsilon_0-2\epsilon_1,\forall\nu\in A_{\epsilon_1}. \tag{26}$$

Step 3. Finally, we show that f_0 is the required Lipschitz function. For each $\mu\in A$ there exists an $\nu_\mu\in A_{\epsilon_1}$ such that $W_1(\mu,\nu_\mu)<\epsilon_1$ and hence $\int fd\mu\geq\int fd\nu_\mu-\epsilon_1,\forall f\in[\mathscr{L}]_1$. This implies,

$$\int f_0d\mu\geq\int f_0d\nu_\mu-\epsilon_1. \tag{27}$$

Similarly, $W_1(\mu_0,\nu_0)<\epsilon_1$ implies

$$\int f_0d\nu_0\geq\int f_0d\mu_0-\epsilon_1. \tag{28}$$

Combining (26) – (28), we obtain $\int f_0 d\mu \geq \int f_0 d\mu_0 + \epsilon_0 - 4\epsilon_1 = \int f_0 d\mu_0 + \frac{\epsilon_0}{5}, \forall \mu \in A$, which yields the required separability. □

Analogously, we can obtain the same separation result in the space of $\mathscr{P}_{W_1} \times \mathbb{R}$ which will be used in the next subsection.

THEOREM B.11. *Let $\tilde{A}$ be a convex and closed subset in $\mathscr{P}_{W_1} \times \mathbb{R}$ with the metric*

$$\tilde{d}((\mu, r_1), (\nu, r_2)) := W_1(\mu, \nu) \vee |r_1 - r_2|.$$

Let $(\mu_0, r_0) \in \mathscr{P}_{W_1} \times \mathbb{R}$ be a point not contained in A. Then there exists a function $(f, \alpha) \in \mathscr{L} \in \mathbb{R}$ such that

$$\int f d\mu + \alpha r > \int f d\mu_0 + \alpha r_0, \forall (\mu, r) \in \tilde{A}.$$

PROOF. The proof follows similarly as the proof of Theorem B.10 and is therefore omitted. □

B.5 Proof of the duality

We shall prove the main duality theorem in this subsection. It is worth to be mentioned that our proofs of Lemma B.14 and the duality theorem (Theorem B.15) mainly follow the line of the proof of Theorem 2.2.15 of Ioan-Bot et al. (2009).

DEFINITION B.12. *For a function $\phi : \mathscr{P}_{W_1} \to \bar{\mathbb{R}}$, its domain is $dom(\phi) := \{\mu \in \mathscr{P}_{W_1} : \phi(\mu) < +\infty\}$. ϕ is said to be* proper *if $\phi(\mu) > -\infty$ for all $\mu \in \mathscr{P}_{W_1}$ and $dom(\phi) \neq \emptyset$. The epigraph of ϕ is defined as $epi(\phi) := \{(\mu, r) \in \mathscr{P}_{W_1} \times \mathbb{R} \mid \phi(\mu) \leq r\}$.*

PROPOSITION B.13. *If $\phi : \mathscr{P}_{W_1} \to \bar{\mathbb{R}}$ is convex and lower semicontinuous (equipped with the metric W_1), then $epi(\phi)$ is convex and closed w.r.t. the metric $\tilde{d}((\mu, a), (\nu, b)) := W_1(\mu, \nu) \vee |a - b|$.*

PROOF. The convexity is easy to verify. We show here only the closeness. Let $\{(\mu_n, a_n) \in \mathscr{P}_{W_1} \times \mathbb{R}, n = 1, 2, \ldots\}$ be any converging sequence in epi(ϕ) under the metric $\tilde{d}$. Hence, we have $\phi(\mu_n) \leq a_n, \forall n = 1, 2, \ldots$. Note that converging under $\tilde{d}$ implies that μ_n converges to a probability measure $\mu_0 \in \mathscr{P}_{W_1}$ (equipped with W_1) and a_n converges to a real number $a_0 \in \mathbb{R}$. By the lower semicontinuity, we have $\phi(\mu_0) \leq \liminf_{n\to\infty} \phi(\mu_n) \leq \liminf_{n\to\infty} a_n = a_0$, which implies that $(\mu_0, a_0) \in$ epi(ϕ). Hence, epi(ϕ) is closed. □

We define the following set

$$M := \left\{ (f, \eta) \in \mathscr{L} \times \mathbb{R} \middle| \int f d\mu + \eta \leq \phi(\mu), \forall \mu \in \mathscr{P}_{W_1} \right\}. \tag{29}$$

LEMMA B.14. *Let $\phi : \mathscr{P}_{W_1} \to \bar{\mathbb{R}}$ be a proper, convex and lower semicontinuous function (equipped with the metric W_1). Then the set M is not empty.*

PROOF. Since ϕ is proper, there exists an element $\nu \in \mathscr{P}_{W_1}$ such that $\phi(\nu) \in \mathbb{R}$. Then epi$(\phi) \neq \emptyset$ and $(\nu, \phi(\nu) - 1) \notin$ epi(ϕ). By Theorem B.11, there exists $(f_0, \eta_0) \in \mathscr{L} \times \mathbb{R}$ such that

$$\int f_0 d\nu + \eta_0(\phi(\nu) - 1) < \int f_0 d\mu + \eta_0 r, \forall (\mu, r) \in \text{epi}(f).$$

Since $(\nu, \phi(\nu)) \in$ epi(f), we have $\eta_0 > 0$ and $(1/\eta_0)\left(\int f_0 d\nu - \int f_0 d\mu\right) + \phi(\nu) - 1 < r, \forall (\mu, r) \in$ epi(ϕ). For any $\mu \in$ dom(f), $(\mu, \phi(\mu)) \in$ epi(ϕ) and hence

$$(1/\eta_0)\left(\int f_0 d\nu - \int f_0 d\mu\right) + \phi(\nu) - 1 < \phi(\mu).$$

This inequality can in fact be extended to all $\mu \in \mathscr{P}_{W_1}$. Thus, $(-f_0/\eta_0, \phi(\nu) - 1 + \int f_0 d\nu/\eta_0)$ is one element of M. □

It is easy to verify that under the same condition as in the above lemma, we have the following equation for the second dual function ϕ^c:

$$\phi^c(\mu_0) = \sup_{(f,\eta)\in M} \left\{ \int f d\mu_0 + \eta \right\}.$$

We restate Theorem 4.1 as follows.

THEOREM B.15. *Let $\phi : \mathscr{P}_{W_1} \to \bar{\mathbb{R}}$ be a proper, convex and lower semicontinuous function on $(\mathscr{P}_{W_1}, W_1)$. Then $\phi(\mu) = \phi^c(\mu)$, $\forall \mu \in \mathscr{P}_{W_1}$.*

PROOF. By definition, we have $\rho(f) \geq \int f d\mu - \phi(\mu), \forall (\mu, f) \in \mathscr{P}_{W_1} \times \mathscr{L}$, which implies $\int f d\mu - \rho(f) \leq \phi(\mu), \forall (\mu, f) \in \mathscr{P}_{W_1} \times \mathscr{L}$, and hence $\phi^c(\mu) \leq \phi(\mu), \forall \mu \in \mathscr{P}_{W_1}$.

Below we show that $\phi(\mu) \leq \phi^c(\mu), \forall \mu \in \mathscr{P}_{W_1}$. Assume the opposite: there exists a $\mu_0 \in \mathscr{P}_{W_1}$ and $r_0 \in \mathbb{R}$ such that

$$\phi(\mu_0) > r_0 > \phi^c(\mu_0) = \sup_{(f,\eta) \in M} \left\{ \int f d\mu_0 + \eta \right\}, \tag{30}$$

where M is defined in (29). It is clear that $(\mu_0, r_0) \notin \text{epi}(\phi)$. Note that by Proposition B.13, $\text{epi}(\phi)$ is closed. Hence, by Theorem B.11, there exists a $(f_0, \eta_0) \in \mathscr{L} \times \mathbb{R}$ such that

$$\int f_0 d\mu + \eta_0 r > \int f_0 d\mu_0 + \eta_0 r_0 + \epsilon, \forall (\mu, r) \in \text{epi}(\phi), \tag{31}$$

Note that if $(\mu, r) \in \text{epi}(\phi)$, then $(\mu, r + s) \in A$ for any $s \geq 0$. Thus $\eta_0 \geq 0$.

Suppose $\phi(\mu_0) \in \mathbb{R}$. Using the fact that $(\mu_0, \phi(\mu_0)) \in A$, we obtain $\eta_0(\phi(\mu_0) - r_0) > 0$. Hence $\eta_0 > 0$. Since $(\mu, \phi(\mu)) \in \mathbb{R}$, we obtain $\phi(\mu) > \frac{1}{\eta_0} \int f_0 d\mu_0 - \frac{1}{\eta_0} \int f_0 d\mu + r_0$. Now setting $\eta = r_0 + \frac{1}{\eta_0} \int f_0 d\mu_0$ and $f = -\frac{f_0}{\eta_0}$ in (30), we have $\int f d\mu_0 + \eta = r_0$, which contradicts (30).

Hence, $\phi(\mu_0) = +\infty$. If $\eta_0 > 0$, the contradiction remains. Thus $\eta_0 = 0$ and (31) becomes $\int f_0 d\mu > \int f_0 d\mu_0 + \epsilon, \forall (\mu, r) \in \text{epi}(\phi)$. By Lemma B.14, the set M defined in (29) is not empty. Hence, there exists $(f_1, a_1) \in M$ such that $\int f_1 d\mu + a_1 \leq \phi(\mu), \forall \mu \in \mathscr{P}_{W_1}$. By the assumption made in (30), $b := (r_0 - \int f_1 d\mu_0 - a_1)/\epsilon > 0$. For all $\mu \in \text{dom}(f)$,

$$\begin{aligned} &\int (f_1 - b f_0) d\mu + \int b f_0 d\mu_0 + a_1 + b\epsilon \\ &= \int f_1 d\mu + a_1 + b \left(- \int f_0 d\mu + \int f_0 d\mu_0 + \epsilon \right) \leq \int f_1 d\mu + a_1 \leq \phi(\mu). \end{aligned} \tag{32}$$

This can be extended to all $\mu \in \mathscr{P}_{W_1}$. Hence, $(f_1 - b f_0, \int b f_0 d\mu_0 + a_1 + b\epsilon) \in M$. Taking $\mu = \mu_0$ in (32), we have $\int (f_1 - b f_0) d\mu_0 + \int b f_0 d\mu_0 + a_1 + b\epsilon = r_0$, which again contradicts (30). □

References

Alan Bain, Dan Crisan, D Dawson, D Geman, G Grimmett, I Karatzas, F Kelly, Y Le Jan, B Øksendal, G Papanicolaou, et al. *Fundamentals of stochastic filtering*, volume 3. Springer, 2009.

D.P. Bertsekas and S.E. Shreve. *Stochastic Optimal Control: The Discrete Time Case*. Academic Press, 1978.

J.M. Borwein and A.S. Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Springer, 2nd edition, 2006.

S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

E.A. Feinberg, P.O. Kasyanov, and M.Z. Zgurovsky. Partially observable total-cost Markov decision processes with weakly continuous transition probabilities. *Mathematics of Operations Research*, 2016.

H. Föllmer and A. Schied. *Stochastic Finance*. Walter de Gruyter & Co., Berlin, 2004. Extended edition.

O. Hernández-Lerma. *Adaptive Markov Control Processes*. Springer, 1989.

O. Hernández-Lerma and J.B. Lasserre. *Discrete-time Markov Control Processes: Basic Optimality Criteria*. Springer, 1996.

O. Hernández-Lerma and J.B. Lasserre. *Further Topics on Discrete-Time Markov Control Processes*. Springer Verlag, 1999.

R. Ioan-Bot, S.-M. Grad, and G. Wanka. *Duality in Vector Optimization*. Springer Science & Business Media, 2009.

Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1):99–134, 1998.

W.S. Lovejoy. A survey of algorithmic methods for partially observed Markov decision processes. *Annals of Operations Research*, 28(1):47–65, 1991.

D.G. Luenberger. *Optimization by Vector Space Methods*. JohnWiley & Sons Inc., New York, 1969.

J. Pineau, G. Gordon, and S. Thrun. Anytime point-based approximations for large pomdps. *Journal of Artificial Intelligence Research*, pages 335–380, 2006.

R.P.N. Rao. Decision making under uncertainty: a neural model based on partially observable Markov decision processes. *Frontiers in computational neuroscience*, 4, 2010.

H.L. Royden. *Real Analysis*. Macmillan, 2nd edition, 1968.

Y. Sawaragi and T. Yoshikawa. Discrete-time Markovian decision processes with incomplete state observation. *The Annals of Mathematical Statistics*, 41(1):78–86, 1970.

G. Shani, J. Pineau, and R. Kaplow. A survey of point-based POMDP solvers. *Autonomous Agents and Multi-Agent Systems*, 27(1):1–51, 2013.

R.D. Smallwood and E.J. Sondik. The optimal control of partially observable Markov processes over a finite horizon. *Operations Research*, 21(5):1071–1088, 1973.

E.J. Sondik. The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs. *Operations Research*, 26(2):282–304, 1978.

C. Villani. *Optimal Transport*. Springer, 2009.

N. Weaver. *Lipschitz Algebras*. World Scientific, 1999.

C. Zălinescu. *Convex Analysis in General Vector Spaces*. World Scientific, 2002.